

---

# Ensemble Validation: Selectivity has a Price, but Variety is Free

---

Eric Bax\* and Farshad Kooti†

## Abstract

If classifiers are selected from a hypothesis class to form an ensemble, bounds on average error rate over the selected classifiers include a component for selectivity, which grows as the fraction of hypothesis classifiers selected for the ensemble shrinks, and a component for variety, which grows with the size of the hypothesis class or in-sample data set. We show that the component for selectivity asymptotically dominates the component for variety, meaning that variety is essentially free.

## 1 Introduction

In machine learning, ensemble methods combine multiple classifiers into a single classifier. Ensemble methods include bagging [1], boosting [2, 3], forests (of decision trees) [4, 5, 6], and stacking [7]. In this paper, we focus on perhaps the simplest ensemble classifier: the Gibbs classifier. To classify an example, a Gibbs classifier chooses one of the classifiers from its ensemble at random and applies that classifier to the example. (The Gibbs classifiers considered in this paper use a uniform distribution to select a classifier from the ensemble.) Our goal is to develop out-of-sample error bounds for Gibbs classifiers. The error rate of a Gibbs classifier is the average error rate over the classifiers its ensemble.

The PAC-Bayes technique [8, 9, 10] bounds the average out-of-sample error rate over a set of classifiers. It applies to the following setup: start with a prior distribution over a set of hypothesis classifiers, use a set of in-sample data to select a posterior distribution over classifiers, and use that posterior as the basis for an ensemble classifier. PAC-Bayes error bounds include a term for selectivity that grows as the Kullback-Leibler (KL) divergence [11] (or another divergence [10]) between the posterior and prior grows and a term that grows with the number of in-sample examples.

Our setup is similar but slightly different: we begin with a set of hypothesis classifiers and a set of validation data for each classifier, choose a number of classifiers, then use the validation data to select that number of classifiers from the hypothesis set for the ensemble. For each classifier, the validation data must be independent of any training data for the classifier. However, the validation data for one hypothesis classifier may be used to train another hypothesis classifier. Our statistics over classifiers are a restriction of the PAC-Bayes setup: we use a uniform prior over the hypothesis set and a uniform posterior over the ensemble classifiers. PAC-Bayes uses a restricted version of our data assumptions: for PAC-Bayes all in-sample data are the validation data for every hypothesis classifier.

We use nearly uniform error bounds [12] in this paper. Nearly uniform error bounds on error rates over a set of classifiers allow the bounds for some classifiers to fail in exchange for tighter bounds on the other classifiers. This idea has been applied to strengthen error bounds over hypothesis classes that contain duplicate classifiers or consist of subsets of classifiers with high rates of agreement. For those cases, if the number of bound failures allowed is less than the number of duplicates (or near

---

\*Yahoo, baxhome@yahoo.com

†Facebook

duplicates) for the classifier selected by training, then the nearly uniform bound applies to at least one of the duplicates (or near duplicates.)

In this paper, we use nearly uniform error bounds over the set of hypothesis classifiers to infer an average error bound over the ensemble of selected classifiers. Allowing a number of hypothesis classifier bound failures that is a small fraction of the number of ensemble classifiers only increases the inferred ensemble average bound by that small fraction, due to averaging. In exchange, we get tighter bounds on the remaining hypothesis classifiers and therefore on the remaining ensemble classifiers. This decreases the ensemble average error bound. We derive the results in this paper by selecting the number of allowed failures to mediate this tradeoff.

Section 2 briefly reviews uniform and nearly uniform error bounds. Section 3 shows how to infer ensemble average error bounds from uniform and nearly uniform error bounds. Section 4 shows how to expand nearly uniform error bounds into telescoping error bounds. Section 5 shows how to select parameters for telescoping error bounds to make the price of variety asymptotically negligible compared to the price of selectivity. Section 6 concludes with a discussion of how to apply the error bounds we derive.

## 2 Uniform and Nearly Uniform Validation

Let's begin with a brief review of uniform validation. Let  $m$  be the number of classifiers in the hypothesis class. Let  $n$  be the number of validation examples, with known inputs and labels, available for each classifier in the hypothesis class. Assume the validation examples for each classifier are drawn independently of any data used to choose the classifier for the class and are drawn i.i.d. from an unknown input-label distribution  $D$ . (The validation data may be the same, partially overlap, or not overlap at all for different classifiers in the class.) For each classifier  $i$  in the class, let  $p_i$  be the error rate of classifier  $i$  over its validation data. Let  $p_i^*$  be the (unknown) error rate of classifier  $i$  over  $D$ . Call  $p_i^*$  the actual error rate or out-of-sample error rate of classifier  $i$ .

Applying Hoeffding bounds [13], for each classifier  $i$ , for  $\delta > 0$ :

$$\Pr \left\{ p_i^* \geq p_i + \sqrt{\frac{\ln \frac{1}{\delta}}{2n}} \right\} \leq \delta.$$

Substitute  $\frac{\delta}{m}$  for  $\delta$ :

$$\Pr \left\{ p_i^* \geq p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right\} \leq \frac{\delta}{m}.$$

Use the sum bound for the probability of a union to derive a uniform bound:

$$\Pr \left\{ \exists i : p_i^* \geq p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right\} \leq \sum_{i=1}^m \Pr \left\{ p_i^* \geq p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right\} \leq m \frac{\delta}{m} = \delta. \quad (1)$$

This is the standard PAC error bound for uniform convergence of classifier validation error rates to their actual error rates, from Vapnik and Chervonenkis (VC) [14].

Uniform validation requires that every classifier in the class has a successful validation in the sense that its validation error rate is close to its actual error rate. In contrast, nearly uniform validation allows some classifiers to have validation error rates far from their actual error rates. Call these misvalidations.

To derive nearly uniform error bounds, select a number  $j$  of misvalidations to allow. The probability of more than  $j$  misvalidations is maximized if there are exactly  $j + 1$  misvalidations in each case of any misvalidations. For  $j = 0$ , this is just the sum bound for the union of probabilities: the probability of at least one misvalidation is maximized if there is exactly one misvalidation when there is any misvalidation – any more and the probability of misvalidation will be less than the sum of individual classifier misvalidation probabilities. For  $j = 1$ , which allows a single misvalidation, the greatest probability of two or more misvalidations is achieved if any misvalidation that occurs is accompanied by exactly one more.

So the probability of more than  $j$  misvalidations is at most  $\frac{1}{j+1}$  of the probability of one or more misvalidations. Apply this rule to produce a nearly uniform version of Inequality 1:

$$\Pr \left\{ \left| \left\{ i \mid p_i^* \geq p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right\} \right| > j \right\} \leq \frac{\delta}{j+1}.$$

Substitute  $\delta(j+1)$  for  $\delta$ :

$$\Pr \left\{ \left| \left\{ i \mid p_i^* \geq p_i + \sqrt{\frac{\ln \frac{m}{\delta(j+1)}}{2n}} \right\} \right| > j \right\} \leq \delta. \quad (2)$$

This is a nearly uniform PAC error bound. Comparing it to the uniform error bound in Inequality 1, in exchange for allowing  $j$  misvalidations, we can subtract  $\ln(j+1)$  from the numerator in the square root.

### 3 Validation for an Ensemble

Let  $S$  be a subset of the  $m$  hypothesis classifiers. Refer to  $S$  as the selected classifiers or the ensemble. Classifiers may be selected for the ensemble based on their validation error rates. Let  $s = |S|$ .

Our goal is to derive an error bound for the average out-of-sample error rate over the ensemble classifiers. Let  $E_S$  be the expectation over ensemble classifiers. Specifically, let  $E_S p_i$  be the average validation error rate, and let  $E_S p_i^*$  be the average out-of-sample error rate. (Then  $E_S p_i^*$  is the out-of-sample error rate of a Gibbs classifier based on a uniform distribution over the ensemble classifiers.)

For an average bound based on uniform bounds, multiply both sides of the inequality in the probability from Inequality 1 by  $\frac{1}{s}$ :

$$\Pr \left\{ \exists i : \frac{p_i^*}{s} \geq \frac{1}{s} \left( p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right) \right\} \leq \delta.$$

So

$$\Pr \left\{ \sum_{i \in S} \frac{p_i^*}{s} \geq \sum_{i \in S} \frac{1}{s} \left( p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right) \right\} \leq \delta,$$

and

$$\Pr \left\{ E_S p_i^* \geq E_S p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right\} \leq \delta. \quad (3)$$

Now consider an average bound based on nearly uniform bounds. In the worst case (with the nearly uniform bound holding), the  $j$  allowed misvalidations occur, they are all for classifiers in the ensemble, and those classifiers have the least validation error rates in the ensemble and out-of-sample error rates of one. In that case, each misvalidation adds  $\frac{1}{s}$  to the average error bound, minus the contribution it would have made to the bound if not for misvalidation:

$$\frac{1}{s} \left( p_i + \sqrt{\frac{\ln \frac{m}{\delta}}{2n}} \right).$$

Let  $L$  index the  $j$  classifiers in  $S$  with the least validation error rates. Then we have:

$$\Pr \left\{ \sum_{i \in S} \frac{p_i^*}{s} \geq \sum_{i \in S} \frac{1}{s} \left( p_i + \sqrt{\frac{\ln \frac{m}{\delta(j+1)}}{2n}} \right) + \frac{j}{s} - \sum_{i \in L} \left( p_i + \sqrt{\frac{\ln \frac{m}{\delta(j+1)}}{2n}} \right) \right\} \leq \delta.$$

Cancel the terms in square roots to get:

$$\Pr \left\{ E_S p_i^* \geq E_S p_i + \left( 1 - \frac{j}{s} \right) \sqrt{\frac{\ln \frac{m}{\delta(j+1)}}{2n}} + \frac{j}{s} (1 - E_L p_i) \right\} \leq \delta.$$

To produce a bound on the difference between average validation error rate and average out-of-sample error rate before observing the validation error rates, consider that in the worst case  $\forall i \in L : p_i = 0$ , so  $E_L p_i = 0$ :

$$\Pr \left\{ E_S p_i^* - E_S p_i \geq \left(1 - \frac{j}{s}\right) \sqrt{\frac{\ln \frac{m}{\delta(j+1)}}{2n}} + \frac{j}{s} \right\} \leq \delta. \quad (4)$$

Let

$$\hat{\epsilon}(j, \delta) = \min \left( \sqrt{\frac{\ln \frac{m}{\delta(j+1)}}{2n}}, 1 \right).$$

Then we can rewrite Inequality 4 as

$$\Pr \left\{ E_S p_i^* - E_S p_i \geq \left(1 - \frac{j}{s}\right) \hat{\epsilon}(j, \delta) + \frac{j}{s} \right\} \leq \delta. \quad (5)$$

## 4 Telescoping Bounds

In the bound in Inequality 5, we have assumed that misvalidated classifiers have error rate one, regardless of their validation error rates, because one is a trivial upper bound on out-of-sample error rate. Instead, we could use a loose bound for the misvalidations that is tighter than the trivial bound. Since this “backstop” uniform bound must hold simultaneously with our original nearly uniform bound, we have to split  $\delta$  between them and appeal to a sum bound on the probability of either of these bounds failing. Let  $\delta_1 + \delta_2 = \delta$ . Then we have:

$$\Pr \left\{ E_S p_i^* - E_S p_i \geq \left(1 - \frac{j}{s}\right) \hat{\epsilon}(j, \delta_1) + \frac{j}{s} \hat{\epsilon}(0, \delta_2) \right\} \leq \delta,$$

where  $\hat{\epsilon}(0, \delta_2)$  is the backstop uniform bound.

Extending this idea produces “telescoping” bounds: a uniform bound and a sequence of nearly uniform bounds, each one tighter but allowing more misvalidations than the previous one. Each bound’s  $\hat{\epsilon}$  applies to the misvalidations allowed by the next bound but not by the present one. The final bound applies to the classifiers that are not misvalidated. Listing the bounds from right to left (it is math, after all) gives the telescoping bound:

$$\Pr \left\{ E_S p_i^* - E_S p_i \geq \left(1 - \frac{\sum_{i=1}^t j_i}{s}\right) \hat{\epsilon}\left(\sum_{i=1}^t j_i, \delta_1\right) + \frac{j_1}{s} \hat{\epsilon}\left(\sum_{i=2}^t j_i, \delta_2\right) + \dots + \frac{j_{t-1}}{s} \hat{\epsilon}(j_t, \delta_t) + \frac{j_t}{s} \hat{\epsilon}(0, \delta_{t+1}) \right\} \leq \delta_1 + \dots + \delta_{t+1}. \quad (6)$$

We can optimize this bound by selecting  $j_1, \dots, j_t$  and  $\delta_1, \dots, \delta_{t+1}$  through dynamic programming. For details, refer to Appendix A.

## 5 Selectivity and Variety

We will show that we can use telescoping bounds to make the price of variety arbitrarily small (asymptotically) compared to the price of selectivity. Let’s begin with the non-telescoping bound from Inequality 5:

$$\Pr \left\{ E_S p_i^* - E_S p_i \geq \left(1 - \frac{j}{s}\right) \hat{\epsilon}(j, \delta) + \frac{j}{s} \right\} \leq \delta.$$

Let  $j = \frac{s}{\sqrt{2n}}$  and use one instead of  $\left(1 - \frac{j}{s}\right)$  for simplicity to get a bound:

$$\Pr \{ E_S p_i^* - E_S p_i \geq \epsilon_0 \} \leq \delta,$$

where

$$\epsilon_0 = \frac{1}{\sqrt{2n}} \sqrt{\ln \frac{m}{s} + \ln \sqrt{2n} + \ln \frac{1}{\delta}} + \frac{1}{\sqrt{2n}}.$$

Except for the  $\frac{1}{\sqrt{2n}}$  term on the far right, this is similar to a PAC-Bayes result from McAllester [8]. (The KL-divergence [11] from McAllester's bound is  $\ln \frac{m}{s}$  for us, using a uniform prior over the hypothesis class and the uniform posterior over the ensemble.) We can interpret the term  $\ln \frac{m}{s}$  as the price for selectivity; it increases as we shrink the ensemble. We can interpret the term  $\ln \sqrt{2n}$  as a price for variety.

The price for selectivity makes sense: for a single classifier ( $s = 1$ ), it gives the  $\ln m$  term from the standard VC error bounds [14] for selecting one classifier from a size- $m$  hypothesis class (our Inequality 1). For an average over all classifiers in the hypothesis class ( $s = m$ ), it gives  $\ln \frac{m}{m} = \ln 1 = 0$ , which is consistent with a VC error bound for a hypothesis class consisting of a single hypothesis – the Gibbs classifier composed of all hypothesis classifiers. (For more details refer to [15].)

The price for variety seems like it could be an artifact of the derivation of the bound – something we can remove or at least reduce. To do that, we use telescoping nearly uniform bounds. Consider using a telescoping bound with one nearly uniform bound and a uniform backstop bound. From Inequality 4,

$$\Pr \{E_{Sp_i^*} - E_{Sp_i} \geq \epsilon_1\} \leq \delta_1 + \delta_2,$$

where

$$\epsilon_1 = \left(1 - \frac{j}{s}\right) \hat{\epsilon}(j, \delta_1) + \frac{j}{s} \hat{\epsilon}(0, \delta_2).$$

Let  $j = \frac{s}{n^a}$  and  $\delta_1 = \delta_2 = \frac{\delta}{2}$ . Then

$$\epsilon_1 \leq \frac{1}{\sqrt{2n}} \sqrt{\ln \frac{m}{s} + \ln 2n^a + \ln \frac{1}{\delta}} + \frac{1}{\sqrt{2n^{1+2a}}} \sqrt{\ln 2m + \ln \frac{1}{\delta}}. \quad (7)$$

For  $a > 0$ , and assuming  $m$  does not grow exponentially in  $n$ , the first term of the RHS dominates the second as  $n$  grows. Compare the first term in Inequality 7 for  $\epsilon_1$  to the first term in  $\epsilon_0$ . Notice that  $\ln 2n^a$  in the inequality for  $\epsilon_1$  replaces  $\ln \sqrt{2n}$  in  $\epsilon_0$ . By setting  $a$  close to zero, we can make  $\ln 2n^a$  small compared to  $\ln \sqrt{2n}$ , shrinking the price of variety.

But suppose that  $m$  is exponential in  $n$ . Then the second term in the RHS of Inequality 7 dominates the first, because then  $\ln 2m$  is  $O(n)$ . (Assume  $a < \frac{1}{2}$ .) So, for hypothesis class sizes that grow exponentially with the number of validation samples, the price of variety asymptotically dominates the price of selectivity. To handle these cases, use another layer in the telescoping bound.

Use the bound:

$$\Pr \{E_{Sp_i^*} - E_{Sp_i} \geq \epsilon_2\} \leq \delta_1 + \delta_2,$$

where

$$\epsilon_2 = \left(1 - \frac{j_1 + j_2}{s}\right) \hat{\epsilon}(j_1 + j_2, \delta_1) + \frac{j_1}{s} \hat{\epsilon}(j_2, \delta_2) + \frac{j_2}{s} \hat{\epsilon}(0, 0).$$

Note that  $\hat{\epsilon}(0, 0) = 1$ . Let  $\delta_1 = \delta_2 = \frac{\delta}{2}$ ,  $j_1 = \frac{s}{n^a}$  with  $a < \frac{1}{2}$ , and  $j_2 = \frac{s}{n^b}$  with  $b \geq \frac{1}{2}$ . Then

$$\epsilon_2 \leq \frac{1}{\sqrt{2n}} \sqrt{\ln \frac{m}{s} + \ln 2n^a + \ln \frac{1}{\delta}} + \frac{1}{\sqrt{2n^{1+2a}}} \sqrt{\ln \frac{m}{s} + \ln 2n^b + \ln \frac{1}{\delta}} + \frac{1}{n^b}. \quad (8)$$

For example, set  $a = 0.1$  and  $b = 0.6$ . Then we have:

$$\epsilon_2 \leq \frac{1}{\sqrt{2n}} \sqrt{\ln \frac{m}{s} + 0.1 \ln 2n + \ln \frac{1}{\delta}} + \frac{1}{\sqrt{2n^{1.2}}} \sqrt{\ln \frac{m}{s} + 0.6 \ln 2n + \ln \frac{1}{\delta}} + \frac{1}{n^{0.6}}.$$

Compare Inequality 8 to Inequality 7. Like in Inequality 8, the second term in Inequality 7 is asymptotically dominated by the first if  $a > 0$ . However,  $m$  does not appear in Inequality 8, except in  $\frac{m}{s}$ . So the price of variety is dominated asymptotically by the price of selectivity, even for hypothesis class sizes  $m$  that grow arbitrarily large as a function of the number of validation examples,  $n$ . Since we can set  $a$  arbitrarily close to zero and set  $b > \frac{1}{2}$ , we can make the first term dominant, and we can make  $\ln 2n^a$  (the price of variety in the dominant term) arbitrarily small.

## 6 Discussion

The error bounds developed in this paper allow the validation data for one classifier to be used to choose other classifiers for the hypothesis class. So we can apply these bounds to a Gibbs classifier based on classifiers trained on different splits of the same in-sample data set, with each classifier's validation data the data withheld from its training. Then we can extend an error bound for this holdout Gibbs classifier to an error bound for the single "full" classifier based on all in-sample examples, if we can bound the rate of disagreement between the holdout Gibbs classifier and the full classifier, since the error rate of the full classifier is at most the error rate of the holdout Gibbs classifier plus the rate of disagreement. (For more details, see [16].)

For local classifiers, like nearest neighbor classifiers, we can use statistics to bound the rate of disagreement. For example, suppose we split the in-sample data into 10 subsets, use each subset as the validation data for a one-nearest neighbor classifier consisting of the other nine subsets, use those ten classifiers in a Gibbs classifier, and use the results in this paper to compute an error bound for the Gibbs classifier. To get an error bound for the one-nearest neighbor classifier based on all in-sample examples, we also need to bound the rate of disagreement between that classifier and the Gibbs classifier. The rate of disagreement is at most 10%, because with probability 90%, the Gibbs classifier selects one of the nine classifiers that includes the nearest neighbor to the example being classified. (In the other 10% of cases, the Gibbs classifier selects the one classifier with the nearest neighbor in its withheld validation data.)

For the Gibbs classifier based on only  $m = 10$  classifiers, we should use the average bound based on uniform validation, from Inequality 3. Alternatively, to reduce variance, we could use a Gibbs classifier based on all subsets of 90% of the data. It would also have a 10% bound on its rate of disagreement with the full classifier, because this Gibbs classifier also has a 10% probability of selecting a classifier that does not include the nearest neighbor from the full classifier. For this Gibbs classifier, it is important to use error bounds that have a low cost for variety, because there are  $m = \binom{10n}{n}$  classifiers in the hypothesis set, where  $10n$  is the total number of in-sample examples and  $n$  is the size of each withheld validation set. So we should use the bound from Inequality 8. (The cost for selectivity is zero, since  $s = m$ .) For an efficient method to compute the average validation error rate for this Gibbs classifier, see Appendix B.

For non-local classifiers, where removal of examples can have non-local effects on classification, there are a few options to bound rates of disagreement between holdout Gibbs classifiers and full classifiers, depending on data availability. If the out-of-sample example inputs are available, then compute the rate of disagreement directly over those examples. Otherwise, if there are unlabeled examples available from the out-of-sample input space, then bound the rate of disagreement using standard error bounds, counting disagreements instead of errors.

In the future, it would be interesting to determine whether nearly uniform bounds can improve PAC-Bayesian bounds that are based on arbitrary prior and posterior distributions and on change of measure with respect to divergences. A first step would be to show how nearly uniform bounds apply to arbitrary distributions. Then we would need to show that those improvements can be maintained through a change of measure.

## References

- [1] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [2] Robert E. Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990.
- [3] Yoav Freund and Robert Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *System Sciences*, 55(1):119–139, 1997.
- [4] Tin Kam Ho. Random decision forests. *Proceedings of the 3rd International Conference on Document Analysis and Recognition, Montreal, QC*, page 278282, 1995.
- [5] Tin Kam Ho. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8):832–844, 1998.
- [6] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [7] D. Wolpert. Stacked generalization. *Neural networks*, 5(2):241–259, 1992.

- [8] David A. McAllester. Pac-bayesian model averaging. In *In Proceedings of the Twelfth Annual Conference on Computational Learning Theory*, pages 164–170. ACM Press, 1999.
- [9] John Langford, Matthias Seeger, and Nimrod Megiddo. An improved predictive accuracy bound for averaging classifiers. In *In Proceeding of the Eighteenth International Conference on Machine Learning*, pages 290–297, 2001.
- [10] Luc Begin, Pascal germain, Francois Laviolette, and Jean-Francis Roy. Pac-bayesian bounds based on the renyi divergence. *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.
- [11] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [12] E. Bax and A. Callejas. An error bound based on a worst likely assignment. *Journal of Machine Learning Research*, 9:859–891, 2008.
- [13] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [14] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.
- [15] E. Bax. Validation of average error rate over classifiers. *Pattern Recognition Letters*, pages 127–132, 1998.
- [16] Eric Bax and Ya Le. Some theory for practical classifier validation. *Baylearn*, 2015.

## A How to Optimize Telescoping Bounds

We can use dynamic programming to optimize (over a discrete set of candidate parameter values) the telescoping bound from Inequality 6:

$$\Pr \left\{ E_{Sp_i^*} - E_{Sp_i} \geq \left( 1 - \frac{\sum_{i=1}^t j_i}{s} \right) \hat{\epsilon} \left( \sum_{i=1}^t j_i, \delta_1 \right) + \frac{j_1}{s} \hat{\epsilon} \left( \sum_{i=2}^t j_i, \delta_2 \right) + \dots + \frac{j_{t-1}}{s} \hat{\epsilon} (j_t, \delta_t) + \frac{j_t}{s} \hat{\epsilon} (0, \delta_{t+1}) \right\} \leq \delta_1 + \dots + \delta_{t+1}.$$

Let

$$\epsilon((j_1, \dots, j_t), (\delta_1, \dots, \delta_{t+1})) = \left( 1 - \frac{\sum_{i=1}^t j_i}{s} \right) \hat{\epsilon} \left( \sum_{i=1}^t j_i, \delta_1 \right) + \frac{j_1}{s} \hat{\epsilon} \left( \sum_{i=2}^t j_i, \delta_2 \right) + \dots + \frac{j_{t-1}}{s} \hat{\epsilon} (j_t, \delta_t) + \frac{j_t}{s} \hat{\epsilon} (0, \delta_{t+1}). \quad (9)$$

To optimize the bound, we want to compute:

$$\epsilon^*(\delta) = \min_{\forall i: j_i \in \{0, \dots, s\}, \sum_{i=1}^t j_i \leq s, \sum_{i=1}^{t+1} \delta_i = \delta} \epsilon((j_1, \dots, j_t), (\delta_1, \dots, \delta_{t+1})).$$

We will begin with the terms on the right of the RHS of Equality 9 and work to the left. Let

$$v(i, \Sigma_j, \Sigma_\delta) = \min_{j_i + \dots + j_t = \Sigma_j, \delta_{i+1} + \dots + \delta_{t+1} = \Sigma_\delta} \left[ \frac{j_i}{s} \hat{\epsilon}(j_{i+1} + \dots + j_t, \delta_{i+1}) + \frac{j_{i+1}}{s} \hat{\epsilon}(j_{i+2} + \dots + j_t, \delta_{i+2}) + \dots + \frac{j_t}{s} \hat{\epsilon}(0, \delta_t) \right].$$

The base cases are:

$$v(t, \Sigma_j, \Sigma_\delta) = \frac{\Sigma_j}{s} \hat{\epsilon}(0, \Sigma_\delta).$$

The general recurrence is:

$$v(i, \Sigma_j, \Sigma_\delta) = \min_{j_i \in \{0, \dots, \Sigma_j\}, \delta_{i+1} \in [0, \Sigma_\delta]} \frac{j_i}{s} \hat{\epsilon}(\Sigma_j - j_i, \delta_{i+1}) + v(i+1, \Sigma_j - j_i, \Sigma_\delta - \delta_{i+1}).$$

(Use a discrete set of candidate values for  $\delta_{i+1}$  in  $[0, \Sigma_\delta]$ , for example increments of 0.0001.) The last step is:

$$\epsilon^*(\delta) = \min_{\Sigma_j \in \{0, \dots, s\}, \delta_1 \in [0, \delta]} \left( 1 - \frac{\Sigma_j}{s} \right) \hat{\epsilon}(\Sigma_j, \delta_1) + v(1, \Sigma_j, \delta - \delta_1).$$

## B How to Compute Average Validation Error Rate over Splits for $k$ -Nearest Neighbors

We can use dynamic programming to compute the average validation error rate over all splits of the in-sample data with  $n$  examples withheld for validation and the remaining examples used to form the classifier. Let  $r + n$  be the total number of in-sample examples. Let  $F$  be the full set of in-sample examples. For each split, let  $H$  be the set of  $n$  holdout examples, and let  $g_H$  be the  $k$ -nearest neighbor classifier based on the remaining  $r$  in-sample examples. Then the average holdout error rate over all splits is:

$$E_{H \subseteq F, |H|=n} E_{(x,y) \in H} I(g_H(x) \neq y),$$

where  $I()$  is the indicator function: one if the argument is true and zero otherwise. Reverse the order of expectations:

$$E_{(x,y) \in F} E_{H \subseteq F, |H|=n, (x,y) \in H} I(g_H(x) \neq y).$$

Convert the inner expectation to a probability of error on  $(x, y)$  over classifiers  $g_H$  that have  $(x, y) \in H$ :

$$E_{(x,y) \in F} \Pr \{g_H(x) \neq y\},$$

where the probability is over  $H \subseteq F$  such that  $|H| = n$  and  $(x, y) \in H$ .

For each in-sample example  $(x, y)$ , we can use dynamic programming to compute the probability of error over classifiers  $g_H$  that have  $(x, y) \in H$ . Let  $a_{i,h,v}$  be the probability that  $h$  of the  $i$  nearest neighbors to  $(x, y)$  in  $F - \{(x, y)\}$  are in  $H$ , and  $v$  of the remaining  $i - h$  nearest neighbors have a different label than  $y$ , so that there are  $v$  incorrect votes. The base case is  $i = 0$ . Let  $a_{0,0,0} = 1$  and  $a_{0,h,v} = 0$  for all other  $h$  and  $v$ .

For the general case, let  $(x_i, y_i)$  be the  $i$ th nearest neighbor to  $(x, y)$  in  $F - \{(x, y)\}$ . Let  $p_H(i, h)$  be the probability that the  $i$ th nearest neighbor to  $(x, y)$  is in  $H$ , given that  $h$  of the  $i - 1$  nearest neighbors are in  $H$  and  $(x, y) \in H$  too. Then

$$p_H(i, h) = \max\left(\frac{n - h - 1}{r + n - i}, 0\right).$$

If  $y_i = y$ , then

$$a_{i,h,v} = a_{i-1,h-1,v} p_H(i, h-1) + a_{i-1,h,v} (1 - p_H(i, h)).$$

If  $y_i \neq y$ , then

$$a_{i,h,v} = a_{i-1,h-1,v} p_H(i, h-1) + a_{i-1,h,v-1} (1 - p_H(i, h)).$$

To get the probability of validation error on  $(x, y)$  over classifiers  $g_H$  with  $(x, y) \in H$ , sum over probabilities of  $k$  examples out of the nearest  $i$  not in  $H$  (hence  $k$  votes), with at least  $\frac{k+1}{2}$  incorrect votes:

$$\Pr \{g_H(x) \neq y\} = \sum_{i=k}^{n+k} \sum_{v=\frac{k+1}{2}}^k a_{i,i-k,v}.$$

So the average validation error rate over the  $\binom{r+n}{n}$  splits is:

$$E_{(x,y) \in F} \sum_{i=k}^{n+k} \sum_{v=\frac{k+1}{2}}^k a_{i,i-k,v}.$$